1

Existence and uniqueness for diffusion processes

1.1 Stochastic processes and filtrations

Let $(Y, \mathcal{M}, P)$ be a probability space and denote points in Y by q. A Borel measurable function $\eta: Y \to R^d$ is called an *R^d-valued random variable* on $(Y, \mathcal{M})$. The *expectation* $E\eta$ of η with respect to the probability measure P is defined by $E\eta = \int_Y \eta \, dP$. The R^d-valued random variable η on $(Y, \mathcal{M}, P)$ induces a probability measure ν on R^d defined by $\nu(B) = P(\eta \in B)$, for $B \in \mathcal{B}_d$, where $\mathcal{B}_d$ is the Borel σ-algebra on R^d; ν is called the *distribution of* η.

A family $\mathcal{M}_t$, $t \geq 0$, of σ-algebras satisfying $\mathcal{M}_{t_1} \subseteq \mathcal{M}_{t_2}$ for $t_1 \leq t_2$ and $\mathcal{M}_t \subseteq \mathcal{M}$ for all $t \geq 0$ is called a *filtration* on $(Y, \mathcal{M})$. If for all $t \geq 0$, $\mathcal{M}_t$ includes all sets of P-measure zero, the filtration is called *complete* on $(Y, \mathcal{M}, P)$. If $\mathcal{M}'_t$ is the smallest σ-algebra containing $\mathcal{M}_t$ and all sets of P-measure zero, then $\mathcal{M}'_t$ is called the *completion* of $\mathcal{M}_t$.

A measurable map $\zeta: [0, \infty) \times Y \to R^d$ is called an *R^d-valued stochastic process* on $(Y, \mathcal{M}, P)$. We frequently suppress the variable $q \in Y$ and write $\zeta(t)$ for $\zeta(t, q)$. The stochastic process $\zeta(t)$ on $(Y, \mathcal{M}, P)$ induces a probability measure ν on $\mathcal{B}([0, \infty), R^d)$, the space of measurable functions from $[0, \infty)$ to R^d with the sup-norm topology, defined by $\nu(B) = P(\zeta(\cdot) \in B)$, for Borel sets B in $\mathcal{B}([0, \infty), R^d)$. The measure ν is called the *distribution* of $\zeta(\cdot)$. The distribution ν is uniquely determined by the finite-dimensional distributions of $\zeta(t)$, that is, by the distributions of $(\zeta(t_1), \zeta(t_2), \ldots, \zeta(t_n))$, for all $0 \leq t_1 < t_2 < \ldots < t_n$ and all $n \geq 1$. If $\zeta(t)$ and $\hat{\zeta}(t)$ are two stochastic processes on $(Y, \mathcal{M}, P)$ and $P(\zeta(t) = \hat{\zeta}(t)) = 1$ for each $t \geq 0$, then, using a monotone class theorem, it is not difficult to show that $\zeta(\cdot)$ and $\hat{\zeta}(\cdot)$ possess a common distribution on $\mathcal{B}([0, \infty), R^d)$; the stochastic processes $\zeta(t)$ and $\hat{\zeta}(t)$ are called *versions* of one another. The stochastic process $\zeta(t)$ is called *progressively measurable* with respect to $\mathcal{M}_t$ if $\zeta(t)$ is $\mathcal{M}_t$-measurable for each $t \geq 0$.

1.2 Conditional expectation, martingales and stopping times

If $\mathcal{G} \subset \mathcal{M}$ is itself a σ-algebra and η is a random variable satisfying $E|\eta| < \infty$, then the conditional expectation $E(\eta|\mathcal{G})$ of η given $\mathcal{G}$ is defined by the following two conditions:

(i) $E(\eta|\mathcal{G})$ is $\mathcal{G}$-measurable;
(ii) $\int_A \eta \, dP = \int_A E(\eta|\mathcal{G}) \, dP$, for all $A \in \mathcal{G}$.

A direct application of the Radon–Nikodym theorem shows that the conditional expectation exists and is unique up to sets of P-measure zero. The basic properties of the conditional expectation are as follows:

1. $E(E(\eta|\mathcal{G})) = E\eta$.
2. $E(\eta_1 + \eta_2|\mathcal{G}) = E(\eta_1|\mathcal{G}) + E(\eta_2|\mathcal{G})$.
3. If $\eta \geq 0$, then $E(\eta|\mathcal{G}) \geq 0$.
4. $|E(\eta|\mathcal{G})| \leq E(|\eta|\,|\mathcal{G})$.
5. $E(\eta|\mathcal{G}) = \eta$ if η is $\mathcal{G}$ measurable.
6. $E(\eta|\mathcal{G}) = E\eta$ if η is independent of $\mathcal{G}$, that is, if $P(\{\eta \in B\} \cap A) = P(\eta \in B)P(A)$ for all $B \in \mathcal{B}_d$ and $A \in \mathcal{G}$.
7. $E(\eta_1\eta_2|\mathcal{G}) = \eta_1 E(\eta_2|\mathcal{G})$ if η_1 is $\mathcal{G}$-measurable.
8. (Bounded convergence theorem for conditional expectation) If $\{\eta_j\}_{j=1}^{\infty}$ is bounded and $\eta = \lim_{j\to\infty} \eta_j$ a.s., then $E(\eta|\mathcal{G}) = \lim_{j\to\infty} E(\eta_j|\mathcal{G})$.
9. (Dominated convergence theorem for conditional expectation) If $\eta = \lim_{j\to\infty} \eta_j$ a.s. and $|\eta_j| \leq Z$ for all $j = 1, 2, \ldots$, where $EZ < \infty$, then $E(\eta|\mathcal{G}) = \lim_{j\to\infty} E(\eta_j|\mathcal{G})$.
10. (Jenson's inequality for conditional expectation) If ϕ is convex and $E|\phi(\eta)| < \infty$, then $\phi(E(\eta|\mathcal{G})) \leq E(\phi(\eta)|\mathcal{G})$.

Of course, all the equalities and inequalities given above are to be taken in the P-almost sure sense.

For $B \in \mathcal{B}_d$, let I_B denote the characteristic function of the set B, that is,

$$I_B(x) = \begin{cases} 1, & \text{if } x \in B \\ 0, & \text{if } x \notin B \end{cases}.$$

We will usually denote $E(I_B(\eta)|\mathcal{G})$ by $P(\eta \in B|\mathcal{G})$. If Z is a random variable, we define the σ-algebra generated by Z as the smallest σ-algebra on Y which contains all sets of the form $\{Z \in B\}$ for $B \in \mathcal{B}_d$. This σ-algebra is denoted by $\sigma(Z)$. We will denote $E(\eta|\sigma(Z))$ and $P(\eta \in B|\sigma(Z))$ by $E(\eta|Z)$ and $P(\eta \in B|Z)$ respectively.

A real-valued progressively measurable stochastic process $\zeta(t)$ is called a *martingale* with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ if $E|\zeta(t)| < \infty$, for all $t \geqslant 0$, and $E(\zeta(t)|\mathcal{M}_s) = \zeta(s)$, for all $0 \leqslant s < t < \infty$. It is called a *submartingale* (*supermartingale*) with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ if $E|\zeta(t)| < \infty$, for all $t \geqslant 0$, and $E(\zeta(t)|\mathcal{M}_s) \geqslant \zeta(s)$ $(E(\zeta(t)|\mathcal{M}_s) \leqslant \zeta(s))$, for all $0 \leqslant s < t < \infty$.

An extended real-valued non-negative random variable $\tau: Y \to [0, \infty]$ is called a *stopping time* with respect to $\mathcal{M}_t$ (or an $\mathcal{M}_t$-stopping time) if $\{\tau \leqslant t\} \in \mathcal{M}_t$, for all $t \geqslant 0$. The following paragraph gives a summary of the basic properties of stopping times.

The *σ-algebra 'up to time τ'* is defined by $\mathcal{M}_\tau = \{A \in \mathcal{M}: A \cap \{\tau \leqslant t\} \in \mathcal{M}_t \text{ for all } t \geqslant 0\}$. Of course, one must verify that $\mathcal{M}_\tau$ as defined is in fact a σ-algebra. If $\tau \equiv t$, then $\mathcal{M}_\tau$ reduces to $\mathcal{M}_t$. Furthermore, if $\sigma \leqslant \tau$ are two stopping times, then $\mathcal{M}_\sigma \subseteq \mathcal{M}_\tau$. Thus, the notation $\mathcal{M}_\tau$ and its designation as the σ-algebra 'up to time τ' are appropriate. If σ and τ are two stopping times, then $\sigma \wedge \tau$ is also a stopping time. Every stopping time τ is $\mathcal{M}_\tau$-measurable. If $\zeta(t)$ is progressively measurable with respect to $\mathcal{M}_t$ and τ is finite-valued, then $\zeta(\tau)$ is $\mathcal{M}_\tau$-measurable.

The principal connection between stopping times and martingales is contained in *Doob's optional sampling theorem*. A version of that theorem which is appropriate for our needs is as follows:

Theorem 2.1 (Doob's optional sampling theorem). *Let $\zeta(t)$ be a martingale (submartingale, supermartingale) and let τ be a stopping time with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$. Then $Z(t) \equiv \zeta(t \wedge \tau)$ is also a martingale (submartingale, supermartingale) with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$.*

The following two martingale inequalities will be useful in the sequel:

Theorem 2.2. *If $\zeta(t)$ is a right continuous submartingale, then*

$$P(\sup_{0 \leqslant s \leqslant t} \zeta(s) \geqslant \lambda) \leqslant \frac{1}{\lambda} E\zeta^+(t),$$

for all $\lambda > 0$ and $t > 0$, where $\zeta^+(t) = \zeta(t) \vee 0$.

Theorem 2.3. *If $\zeta(t)$ is a right continuous martingale or a non-negative, right continuous submartingale, then for any $\alpha > 1$,*

$$E[\sup_{0 \leqslant s \leqslant t} |\zeta(s)|^\alpha] \leqslant \left(\frac{\alpha}{\alpha - 1}\right)^\alpha E|\zeta(t)|^\alpha.$$

In Chapter 9 we will need the *martingale convergence* theorem.

Theorem 2.4. *Let* $\zeta(t)$ *be a submartingale with respect to* $(Y, \mathcal{M}, \mathcal{M}_t, P)$. *Assume that* $E|\zeta(t)|$ *is bounded in* t. *Then* $\lim_{t\to\infty} \zeta(t)$ exists a.s.

1.3 Markov processes and semigroups

An R^d-valued stochastic process $\zeta(t)$ on $(Y, \mathcal{M}, \mathcal{M}_t, P)$ which satisfies $P(\zeta(t) \in B|\mathcal{M}_s) = P(\zeta(t) \in B|\zeta(s))$ for $0 \leq s < t < \infty$ and $B \in \mathcal{B}_d$ is called an R^d-valued *Markov process* on $(Y, \mathcal{M}, \mathcal{M}_t, P)$. Suppose that for each $0 \leq s < t$ and each $x \in R^d$ there exists a probability measure $p(s, x, t, \mathrm{d}y)$ on R^d satisfying

$$\left.\begin{array}{ll}\text{(i)} & p(s, x, t, B) \text{ is measurable as a function of } x \text{ for each } 0 \leq s < t \\ & \text{and } B \in \mathcal{B}_d; \\ \text{(ii)} & p(s, x, t, B) = \int_{R^d} p(s, x, u, \mathrm{d}y) p(u, y, t, B) \text{ for all } s < u < t \\ & \text{and all } B \in \mathcal{B}_d.\end{array}\right\} \tag{3.1}$$

Condition (ii) in (3.1) is known as the *Chapman–Kolmogorov equation*. Suppose, further, that $P(\zeta(t) \in B|\mathcal{M}_s) = P(\zeta(t) \in B|\zeta(s)) = p(s, \zeta(s), t, B)$ for all $0 \leq s < t$ and all $B \in \mathcal{B}_d$. Then $p(s, x, t, \mathrm{d}y)$ is called the *transition probability function* for the Markov process $\zeta(t)$. A Markov process with a transition probability function $p(s, x, t, \mathrm{d}y)$ actually satisfies the following condition:

$$P(\zeta(t + \cdot) \in B|\mathcal{M}_t) = P(\zeta(t + \cdot) \in B|\zeta(t)), \text{ for } t \geq 0$$
$$\text{and Borel sets } B \text{ in } \mathcal{B}([0, \infty), R^d). \tag{3.2}$$

The proof is left as exercise (1.2).

In the special case that $p(s, x, t, \mathrm{d}y) = p(0, x, t - s, \mathrm{d}y)$, for all $0 \leq s < t$ and all $x \in R^d$, we will use the notation $p(t, x, \mathrm{d}y) = p(s, x, s + t, \mathrm{d}y)$. In this case, $P(\zeta(t) \in B|\mathcal{M}_s) = p(t - s, \zeta(s), B)$ depends only on the position of the process $\zeta(s)$ and on the difference $t - s$. Such a Markov process is called *time-homogeneous*. All the Markov processes we consider in this book will be time-homogeneous.

Given a time-homogeneous transition probability function $p(t, x, \mathrm{d}y)$, the Kolmogorov construction (see, for example, [Breiman (1968) or (1992)]) guarantees the existence of a measurable space $(Y, \mathcal{M})$ with filtration $\mathcal{M}_t$ and a progressively measurable stochastic process $\zeta(t)$: $[0, \infty) \times Y \to R^d$ such that for each probability measure μ on R^d, there exists a probability measure P_μ on $(Y, \mathcal{M})$ satisfying

(i) $P_\mu(\zeta(0) \in B) = \mu(B)$ for $B \in \mathcal{B}_d$;
(ii) $P_\mu(\zeta(t) \in B|\mathcal{M}_s) = p(t - s, \zeta(s), B)$, for $0 \leq s < t$ and $B \in \mathcal{B}_d$.

The process $\zeta(t)$ is then a Markov process on $(Y, \mathcal{M}, \mathcal{M}_t, P_\mu)$. Its distribution is uniquely determined by the transition probability function $p(t, x, dy)$ and the initial distribution μ. In the special case where $\mu = \delta_x$ is the atom at some $x \in R^d$, we use the notation P_x. Of course, $P_\mu = \int_{R^d} P_x \mu(dx)$. The expectation corresponding to P_μ is denoted by E_μ.

Let $\mathcal{B}_b(R^d)$ denote the space of real-valued bounded measurable functions on R^d and define the operator $T_t: \mathcal{B}_b(R^d) \to \mathcal{B}_b(R^d)$ by

$$T_t f(x) = E_x f(\zeta(t)) = \int_{R^d} f(y) p(t, x, dy).$$

Since

$$\begin{aligned} T_{t+s} f(x) = E_x f(\zeta(t+s)) &= E_x(E_x f(\zeta(t+s)) | \mathcal{F}_t) \\ &= E_x E_x(f(\zeta(t+s)) | \zeta(t)) \\ &= E_x E_{\zeta(t)} f(\zeta(s)) \\ &= E_x \int_{R^d} f(y) p(s, \zeta(t), dy) \\ &= \int_{R^d} \int_{R^d} f(y) p(s, z, dy) p(t, x, dz), \end{aligned}$$

it follows from the Chapman–Kolmogorov equations that $T_{t+s} = T_t T_s$ for $0 \leqslant s < t$; therefore T_t defines a *semigroup*. Furthermore, letting $\|\cdot\|$ denote the sup-norm on $\mathcal{B}_b(R^d)$, it is clear that T_t is a *contraction semigroup* on $\mathcal{B}_b(R^d)$, that is, $\|T_t f\| \leqslant \|f\|$ for $t \geqslant 0$. (Additional material on semigroups may be found in Chapter 3, Section 3.4 and in exercise 3.3.)

A Markov process $\zeta(t)$ on $(Y, \mathcal{M}, \mathcal{M}_t, P)$ with transition probability function $p(t, x, dy)$ is called a *strong Markov process* if $P(\zeta(\tau + t) \in B | \mathcal{F}_\tau) = p(t, \zeta(\tau), B)$ on $\{\tau < \infty\}$ for all $t \geqslant 0$, all $B \in \mathcal{B}_d$ and all stopping times τ. The Markov process is called a *Feller process* if its semigroup T_t leaves invariant $C_b(R^d)$, the space of bounded continuous functions on R^d. In other words, $\zeta(t)$ is a Feller process if $E_x f(\zeta(t))$ is continuous in $x \in R^d$ for each $t \geqslant 0$ and each $f \in C_b(R^d)$. Using the Markov property, it is not difficult to show that the Feller property is equivalent to the apparently stronger requirement that $x \to E_x \Phi(\zeta(\cdot))$ be continuous for all bounded continuous $\Phi: \mathcal{B}([0, \infty), R^d) \to R$. Thus, in terms of weak continuity of measures (see Chapter 7, Section 3), the Feller property may be defined as the weak continuity of P_x as a function of x. In fact, the Feller property can be shown to be equivalent to a yet stronger condition as follows.

Theorem 3.1. *If $\zeta(t)$ is a Feller process, then $x \to E_x\Phi(\zeta(\cdot))$ is continuous for each bounded $\Phi: \mathcal{B}([0,\infty), R^d) \to R$ which is ν – a.s. continuous, where ν is the distribution of $\zeta(t)$ on $\mathcal{B}([0,\infty), R^d)$.*

1.4 Brownian motion

Consider the transition probability function on R^d defined by $p(t,x,\mathrm{d}y) = p(t,x,y)\,\mathrm{d}y$ with

$$p(t,x,y) = (2\pi t)^{-d/2} \exp\left(\frac{-|x-y|^2}{2t}\right).$$

Then for $t>0$ and $x \in R^d$, $p(t,x,y)$ is the d-dimensional joint Gaussian density with mean zero and covariance tI. The Markov process with this transition probability function and initial distribution μ is called a d-dimensional *Brownian motion* with initial distribution μ. We will denote this process on $(Y, \mathcal{M}, \mathcal{M}_t, P)$ by $B(t) = B(t,q)$. (We suppress the dependence on μ.) An equivalent definition of Brownian motion with initial distribution μ is this:

$B(t)$ is a Markov process on $(Y, \mathcal{M}, \mathcal{M}_t, P)$ satisfying the following three conditions:

(i) $B(0)$ has distribution μ.
(ii) $B(t) - B(s)$ and $B(s)$ are independent for $0 \leq s < t < \infty$.
(iii) $B(t) - B(s)$ has the joint Gaussian distribution with mean zero and covariance $(t-s)I$. (4.1)

As previously noted, any R^d-valued stochastic process induces a measure on $\mathcal{B}([0,\infty), R^d)$ which is called the distribution of the process. We want to show that the distribution of Brownian motion is in fact supported on $C([0,\infty), R^d)$, the space of continuous functions from $[0,\infty)$ to R^d. One must be a bit careful because $C([0,\infty), R^d)$ is not a Borel-measurable subset of $\mathcal{B}([0,\infty), R^d)$. There are many ways to proceed, one of which is as follows. Denote elements of $\mathcal{B}([0,\infty), R^d)$ by ω. Let $\{t_k\}_{k=1}^{\infty}$ be a dense, countable set in $[0,\infty)$ and consider the Borel measurable set

$$\mathcal{U} = \{\omega \in \mathcal{B}([0,\infty), R^d) : \omega \text{ is uniformly continuous on } \{t_k\}_{k=1}^{\infty}\}.$$

One can show that the distribution ν of the Brownian motion satisfies $\nu(\mathcal{U}) = 1$. From this it is easy to show that there exists a continuous version of the Brownian motion. See [Breiman (1968) or (1992)] for a particularly simple derivation or, for example, [Karatzas and Shreve (1988)], where the proof of continuity actually reveals that almost every path is Hölder continuous with exponent γ for any $\gamma < \frac{1}{2}$. We shall always assume that our Brownian motions are continuous. The probability measure on $C([0,\infty), R^d)$ corresponding to Brownian motion in the case where the initial distribution μ is the atom at $x = 0$ is called the (d-dimensional) *Wiener measure*.

Brownian motion is a Feller process and a strong Markov process (exercises 1.3 and 1.4).

We note the following simple lemma, which will be required in Section 7.

Lemma 4.1. *Let $B(t)$ be a Brownian motion on $(Y, \mathcal{M}, \mathcal{M}_t, P)$. Then for each $t \geq 0$, the* σ-algebras $\mathcal{M}_t$ and $\sigma(B(s) - B(t); s \geq t)$ *are independent.*

Proof. By appealing to a standard monotone class theorem, it is enough to show that $P(D|\mathcal{M}_t) = P(D)$, for D of the form $D = \{(B(s_j) - B(t)) \in A_j,\ j = 1, 2, \ldots, n)\}$, where $t < s_1 < s_2 < \ldots < s_n$ and $A_j \in \mathcal{B}_d$, $j = 1, 2, \ldots, n$. Using the Markov property and the fact that $B(t)$ is independent of $B(s) - B(t)$, for $s \geq t$, we have $P(D|\mathcal{M}_t) = P(D|B(t)) = P(D)$. □

1.5 Itô processes

Let $B(t)$ be a d-dimensional Brownian motion on $(Y, \mathcal{M}, \mathcal{M}_t, P)$. For $\theta \in R^d$, define

$$\mathrm{e}_\theta(x) = \exp(i\langle\theta, x\rangle),\ x \in R^d.$$

By property (ii) of (4.1), properties (6) and (7) of the conditional expectation, and the Markov property, we have for $0 \leq t_1 < t_2$,

$$\begin{aligned} E(\mathrm{e}_\theta(B(t_2))|\mathcal{M}_{t_1}) &= E(\exp(\mathrm{i}\langle\theta, B(t_1) + (B(t_2) - B(t_1))\rangle)|\mathcal{M}_{t_1}) \\ &= \mathrm{e}_\theta(B(t_1))E(\mathrm{e}_\theta(B(t_2) - B(t_1))|\mathcal{M}_{t_1}) \\ &= \mathrm{e}_\theta(B(t_1))E(\mathrm{e}_\theta(B(t_2) - B(t_1))|B(t_1)) \\ &= \mathrm{e}_\theta(B(t_1))E\mathrm{e}_\theta(B(t_2) - B(t_1)). \end{aligned}$$

By property (iii) of (4.1),

$$E\mathrm{e}_\theta(B(t_2) - B(t_1)) = \exp\left(\frac{-|\theta|^2}{2}(t_2 - t_1)\right).$$

We conclude that

$$E(\mathrm{e}_\theta(B(t_2))|\mathcal{M}_{t_1}) = \mathrm{e}_\theta(B(t_1))\exp\left(\frac{-|\theta|^2}{2}(t_2-t_1)\right),$$
$$0 \leqslant t_1 < t_2,\ \theta \in R^d. \quad (5.1)$$

Note that (5.1) uniquely determines the distribution of the process $B(t)$ up to its initial distribution (exercise 1.5). Another way of stating (5.1) is that $X_{i\theta}(t)$ is a martingale relative to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ for all $\theta \in R^d$, where $X_{i\theta}(t) = \exp(\mathrm{i}\langle \theta, B(t) - B(0)\rangle + \frac{1}{2}|\theta|^2 t)$.

Now, from (5.1), we have

$$\frac{\mathrm{d}}{\mathrm{d}t}E(\mathrm{e}_\theta(B(t))|\mathcal{M}_{t_1}) = \frac{-|\theta|^2}{2}\mathrm{e}_\theta(B(t_1))\exp\left(\frac{-|\theta|^2}{2}(t-t_1)\right),$$
$$\text{for } 0 \leqslant t_1 < t.$$

Thus, for $t_2 > t_1$,

$$\begin{aligned} &E(\mathrm{e}_\theta(B(t_2)) - \mathrm{e}_\theta(B(t_1))|\mathcal{M}_{t_1}) \\ &\quad = \frac{-|\theta|^2}{2}\int_{t_1}^{t_2}\mathrm{e}_\theta(B(t_1))\exp\left(\frac{-|\theta|^2}{2}(t-t_1)\right)\mathrm{d}t \\ &\quad = \frac{-|\theta|^2}{2}\int_{t_1}^{t_2}\mathrm{E}(\mathrm{e}_\theta(B(t))|\mathcal{M}_{t_1})\,\mathrm{d}t \\ &\quad = E\left(\int_{t_1}^{t_2}\tfrac{1}{2}\Delta\,\mathrm{e}_\theta(B(t))\,\mathrm{d}t|\mathcal{M}_{t_1}\right). \end{aligned}$$

That is,

$$e_\theta(B(t)) - \int_0^t \tfrac{1}{2}\Delta\,\mathrm{e}_\theta(B(s))\,\mathrm{d}s \text{ is a martingale with respect to}$$
$$(Y, \mathcal{M}, \mathcal{M}_t, P), \quad \text{for all } \theta \in R^d. \quad (5.2)$$

Since any $f \in C_0^\infty(R^d)$ can be represented by

$$f(x) = \int_{-\infty}^{\infty}\exp(\mathrm{i}\langle\theta, x\rangle)\phi(\theta)\,\mathrm{d}\theta,$$

for some rapidly decreasing ϕ, it follows easily fom (5.2) and exercise 1.1, that

$$f(B(t)) - \int_0^t \tfrac{1}{2}\Delta f(B(s))\,\mathrm{d}s \text{ is a martingale with respect to}$$
$$(Y, \mathcal{M}, \mathcal{M}_t, P), \quad \text{for all } f \in C_b^2(R^d), \quad (5.3)$$

where $C_b^2(R^d)$ denotes the space of functions on R^d, all of whose partial derivatives up to order 2 are continuous and bounded. Conversely, starting with (5.3), one can obtain (5.1).

Now let

$$\zeta(t) = \sigma B(t) + bt, \tag{5.4}$$

where σ is a constant $d \times d$ matrix and b is a constant d-vector. Letting $a = \sigma\sigma^T$, a calculation similar to the one just carried out reveals that

$$\begin{gathered} X_{i\theta}(t) \text{ is a martingale with respect to } (Y, \mathcal{M}, \mathcal{M}_t, P), \\ \text{where } X_{i\theta}(t) = \exp\left(i\langle \theta, \zeta(t) - \zeta(0) - bt\rangle + \tfrac{1}{2}\langle \theta, a\theta\rangle\right), \end{gathered} \tag{5.5}$$

and that

$$\begin{gathered} f(\zeta(t)) - \int_0^t Lf(\zeta(s))\, \mathrm{d}s \text{ is a martingale with respect to } (Y, \mathcal{M}, \mathcal{M}_t, P), \\ \text{for all } f \in C_b^2(R^d), \end{gathered} \tag{5.6}$$

where

$$L = \tfrac{1}{2}\sum_{i,j=1}^{d} a_{ij}\frac{\partial^2}{\partial x_i \partial x_j} + \sum_{i=1}^{d} b_i \frac{\partial}{\partial x_i}.$$

Furthermore, (5.5) and (5.6) are equivalent.

We now want to construct a Markov process $\zeta(t)$ satisfying (5.6) in the case where

$$L = \tfrac{1}{2}\sum_{i,j=1}^{d} a_{ij}(x)\frac{\partial^2}{\partial x_i \partial x_j} + \sum_{i=1}^{d} b_i(x) \frac{\partial}{\partial x_i}$$

is a variable coefficient elliptic operator. In the light of (5.4), it is tempting to try to arrive at $\zeta(t)$ by solving the 'differential equation'

$$\mathrm{d}\zeta(t) = \sigma(\zeta(t))\,\mathrm{d}B(t) + b(\zeta(t))\,\mathrm{d}t, \tag{5.7}$$

where $\sigma\sigma^{\mathrm{T}}(x) = a(x)$. The problem with this is that $B(t)$ is not of bounded variation; we must develop a suitable integration theory with respect to $\mathrm{d}B(t)$.

It will actually be expedient to develop an integration theory with respect to processes more general than Brownian motion. To this end, let a: $[0, \infty) \times Y \to S_d$ and b: $[0, \infty) \times Y \to R^d$ be progressively measurable on $(Y, \mathcal{M}, \mathcal{M}_t, P)$, where S_d denotes the space of real symmetric non-negative definite $d \times d$ matrices. Let

$$L_t = \tfrac{1}{2}\sum_{i,j=1}^{\mathrm{d}} a_{ij}(t, q)\frac{\partial^2}{\partial x_i \partial x_j} + \sum_{i=1}^{\mathrm{d}} b_i(t, q) \frac{\partial}{\partial x_i}. \tag{5.8}$$

The following theorem generalizes the equivalence of (5.5) and (5.6).

Theorem 5.1. *Let L_t satisfy (5.8) and assume that a_{ij} and b_i are bounded. Let $\zeta(t) = \zeta(t, q)$ be an almost surely continuous process on $(Y, \mathcal{M}, \mathcal{M}_t, P)$.*

(a) *The following four conditions are equivalent:*

(i) *$X_\theta(t)$ is a martingale with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ for all $\theta \in R^d$, where*

$$X_\theta(t) = \exp\left(\left\langle \theta, \zeta(t) - \zeta(0) - \int_0^t b(s)\,ds \right\rangle - \tfrac{1}{2}\int_0^t \langle \theta, a(s)\theta\rangle\,ds\right);$$

(ii) *$X_{i\theta}(t)$ is a martingale with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ for all $\theta \in R^d$, where*

$$X_{i\theta}(t) = \exp\left(i\left\langle \theta, \zeta(t) - \zeta(0) - \int_0^t b(s)\,ds \right\rangle + \tfrac{1}{2}\int_0^t \langle \theta, a(s)\theta\rangle\,ds\right);$$

(iii) *$f(\zeta(t)) - \int_0^t (L_s f)(\zeta(s))\,ds$ is a martingale with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ for all $f \in C_b^2(R^d)$;*

(iv)
$$f(t, \zeta(t)) - \int_0^t \left(\frac{\partial}{\partial s} + L_s\right) f(s, \zeta(s))\,ds$$
is a martingale with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$ for all $f \in C_b^{1,2}([0, \infty) \times R^d)$.

(b) *If $E\exp(\lambda|\zeta(0)|) < \infty$, for all $\lambda \in R$, then $E\exp(\lambda|\zeta(t)|) < \infty$, for all $\lambda \in R$ and all $t \geq 0$. In this case, $f(t, \zeta(t)) - \int_0^t ((\partial/\partial s) + L_s) f(s, \zeta(s))\,ds$ is a martingale for all $f \in C^{1,2}([0, \infty) \times R^d)$ for which there exist constants c_t and k_t for each $t > 0$ such that*

$$\left|\frac{\partial^\alpha f}{\partial x^\alpha}(s, x)\right| \leq c_t \exp(k_t|x|) \text{ and } \left|\frac{\partial f}{\partial s}(s, x)\right| \leq c_t \exp(k_t|x|),$$
for all $x \in R^d$, $0 \leq s \leq t$, $t > 0$ and $|\alpha| \leq 2$.

In particular,

$$\left\langle \theta, \zeta(t) - \int_0^t b(s)\,ds \right\rangle \text{ is a martingale with respect to } (Y, \mathcal{M}, \mathcal{M}_t, P), \tag{5.9}$$

and

$$\left\langle \theta, \zeta(t) - \int_0^t b(s)\,ds \right\rangle^2 - \int_0^t \langle \theta, a(s)\theta\rangle\,ds$$
is a martingale with respect to $(Y, \mathcal{M}, \mathcal{M}_t, P)$. (5.10)

To prove the theorem, we need two results from Chapter 2, so we postpone the proof until Section 5 of Chapter 2.

If a_{ij} and b_i are bounded and $\zeta(t)$ is an almost surely continuous process satisfying the equivalent conditions of Theorem 5.1, we will call $\zeta(t)$ an *Itô process* with *covariance* a and *drift* b. We will write $\zeta(t) \sim I_d(a, b)$.